arXiv:1704.07554v1 [stat.ML] 25 Apr 2017

A relevance-scalability-interpretability tradeoff with temporally evolving user personas

Snigdha Panigrahi* Nadia Fawaz †

April 26, 2017

Abstract

Characterization of consumers has been explored previously in prior works with different goals. The current work presents user personas in a VoD streaming space using a **tenure timeline** and **temporal** behavioral features in the absence of explicit user profiles. A choice of tenure timeline caters to business needs of understanding the evolution and phases of user behavior as their accounts age while temporal characteristics are necessary to capture the dynamic aspect of user personas labels. Our user personas succeed in representing both dominant and niche characterizations and in providing insightful maturation of consumer behavior in the system. The motivation of this work is not limited to discovering latent structure, the emphasis is rather on building interpretable user labels that can lead to immediate and actionable business plans. The two major highlights of our personas are demonstration of *stability* along tenure timelines on a population level, while exhibiting interesting migrations between labels on an individual level. Even on an individual granularity, the personas exhibit stability of dominant characteristics while the niche characterizations are more prone to change. We conclude the paper by illustrating a trade-off between an indispensable trio of guarantees, *relevance-scalability-interpretability*, accomplished by using summary information from personas in CTR predictive models. The proposed approach of uncovering latent temporal structure and consequent insights from evolving personas are broadly applicable to other streaming based products.

Keywords: User personas, Temporal labels, Personalization, CTR prediction, Mixture model

* Statistics Department, Stanford University, CA. Email: snigdha@stanford.edu

†, LinkedIn, CA.

This work was achieved while both authors were employed by Technicolor, Los Altos.

Introduction

Advances in data mining techniques have led to extensive applications [Agrawal et al. (1996); Apte et al. (2002); Chen et al. (2005)] in understanding different aspects of consumer behavior for satisfaction and retention. Consumer segmentation, the idea of dividing a market up into homogeneous segments and targeting each group with a distinct product or message is a basic tool to model similar consumers. It has been explored in diverse sectors- finance Li et al. (2010), health Gummerus et al. (2004), telecommunications Bayer (2010) etc. The current work adopts a *mixture model* approach to construct segments of homogeneous consumers called *user personas* for VoD services from raw transactional logs, using a tenure timeline and temporal behavioral features. These choices of features and comparison along tenures are both extremely integral to explaining consequential temporal, latent structure. Examples of such services in the VoD space include *itune, googleplay, vudu, FandangoNow*, etc; where users pay per piece of content they watch. This is in contrast with subscription based services, where users pay a monthly subscription, such as netflix, hulu plus, amazon video etc. The current work focuses on providing explicit consumer characterizations on- **spending behavior**, **content preference** and **transactional habits**. The main contributions in the psychographic segmentation presented herein are

- aligning user transaction timelines on a **tenure basis** at a **monthly** granularity, a novel choice for a timeline of comparison instead of the usual calendar timeline across users
- constructing temporal feature vectors from transaction logs, that are aggregates of transactions over a month along tenure timeline; such features represent the **evolving behavioral** consumer traits and capture the underlying latent structure
- capturing both the **dominant** and **niche** segments of population and providing interpretable user labels with meaningful characterizations
- capturing stable latent structure on a population level, even as individual profiles keep transforming with age; we inherit the merits of maintaining a **consistent** clustering over time while explaining the changes on a user level accurately
- representing insights on inter-relations between behavioral characteristics like spending behavior and content preferences, genre and recency preferences etc. as layers within user profiles.

Such a construction of temporally evolving user personas with many new perspectives on behavioral characteristics is the "*first of its kind*" in the streaming space, to the best of our knowledge.

While a line of prior works Allenby and Rossi (1998); Besanko et al. (2003); Bhatnagar and Ghose (2004); Wu and Lin (2005); Wedel and Kamakura (2012) has explored characterization of consumers and another independent line of work has contributed to methods of tailoring to users preferences through personalized recommendations Ansari et al. (2000); Sarwar et al. (2001); Schafer et al.

(2007); Pazzani and Billsus (2007); Hu et al. (2008); Ekstrand et al. (2011); Koren et al. (2009), our work leads to a unification of these two important goals to demonstrate the utility of *user personas*. The concluding part of this work illustrates an application of user-persona based personalized recommendations to CTR (Click through Rate) predictions in a VoD transactional dataset. We show that a model based on personas achieves a 3 criteria **relevance - scalability - interpretability** tradeoff, when compared against models that do not use persona information. The CTR model we describe is interesting in its own right as we use persona features in a logistic model trained per item to capture item specific variability. In doing so, we preserve anonymity of individual users as well as privacy of transactions through the use of persona information. We supplement the CTR model with some existing, commonly used collaborative filtering models that can potentially overcome the curse of high dimensionality while still preserving predictive power and business interpretations by using summarized information from personas.

The construction of user personas derives motivation from profiling characterizations that can be translated into business goals and actions; hence, the emphasis is on meaningful labels with clear interpretations. The evolving user labels lend insights into business schemes that can be designed to cater to different consumer segments. As an application to predictive models, we emphasize on substantive gain in computational cost through the use of persona features in the form of soft/hard clustering information and on retaining clarity in meaning of features (as opposed to random projections onto lower dimensional spaces) along side maintenance of predictive ability. This trio of guarantees is extremely important for any practitioner who builds predictive models. More so, the search over feature space for informative predictors is always a costly affair. Persona information can be thereby come handy as natural, informative features, in which case, we retain interpretation of the feature space along with the inherent benefits of easy storage lower dimensional features for future models. Finally, we point out that our methods are not limited to the VoD space and have the scope to be applied to lend similar insights in other product based services.

Related works

Consumer segmentation, driven by the intuition that predictive models of customer behavior based on groups of similar customers outperform a single aggregate model, dates back to Smith (1956), finding basis in Allenby and Rossi (1998), Besanko et al. (2003), Wedel and Kamakura (2012). Predictive models trained per user bucket, obtained through a division of population into homogeneous segments achieve an optimal bias variance trade-off, see Jiang and Tuzhilin (2006). A segmented predictive model can be refined further to an individual level, trained per customer. In doing so, we gain a reduced bias in creating increasingly more "homogeneous" customer groups at the cost of increased variance in estimation as we consider progressively more refined segments containing fewer customers. Thus, there is a classic bias-variance trade-off which

is effectively dealt by integrating customer segmentation into such predictive models, termed as *segmented models* in Jiang and Tuzhilin (2006). We advocate the use of persona features for not just improvement of predictive power, rather as a meaningful, lower dimensional, summary feature space. This can be in turn be used to build models, efficiently implement predictions and can facilitate easy storage for future debugging. The final section explores applications of persona features in CTR predictive models, also discussed in Agarwal et al. (2009); Chapelle et al. (2015).

Various techniques of segmenting consumers have featured in the marketing arena like neural net models in Vellido et al. (1999); Boone and Roehm (2002), latent probabilistic models in Dias and Vermunt (2007); Jedidi et al. (1997), combinatorial optimization based grouping models in Jiang and Tuzhilin (2009). The present work offers a complete mixture model analysis with output as both soft and hard clustering values; thereby, falling in the class of latent paramteric models. Our construction of *user personas* is based on a multinomial mixture model, employing the classic EM algorithm in Dempster et al. (1977) to estimate the mixing proportions and distribution parameters for building user personas. Most part of our raw data consists of count features for which a multinomial model seems a natural choice; except for the spending amounts which we choose to implement the K-means clustering which gives similar results as as the more commonly used parametric Gaussian mixture model Friedman and Russell (1997). Other mixture model implementations find theory and applications in the works Jain et al. (2000), McLachlan and Peel (2004), McLachlan and Krishnan (2007),Vermunt and Magidson (2002), McNicholas and Murphy (2010); Reynolds (2015) etc. In the assessment of clustering performance and determination of number of clusters, we emphasize on clear interpretability of labels in personas; our goal being to render business insights than being limited to discovering latency.

In comparison to the prior art, we lend various new and useful insights on consumer characteristics from the personas constructed in the paper. The personas exhibit stability on a population level even as migrations on an individual level are constantly taking place along the chosen time granularity. This idea of clusters staying faithful to the current data as much as possible while not shifting dramatically from one time-step to the next is also explored in Chakrabarti et al. (2006) and the idea of average network properties reaching an equilibrium can be found in Kossinets and Watts (2006). An immediate merit of such a stability on the population level is that we would not need to learn a completely new way of segmenting data at each time point; such a historically consistent clustering provides greater robustness. We show a divisive structure layered within the personas for each characterization which, unlike hierarchical clustering methods Corpet (1988); Johnson (1967) is not imposed by our approach.

VoD Data set

Dataset considered in this work consists of transaction logs of a subset of $\mathbf{730,000}$ anonymous users from a large-scale streaming VoD service across a

time span of **16** months from January 2014 to April 2015, with over **2** million transactions. Each record in the transaction logs consists of a unique user-id, a unique time-stamp, a unique content-id, the type of transaction–rentals/ purchases, a net price giving the cost of each transaction, and content meta-data such as genres, release year, MPAA ratings corresponding to each transaction. Trouleau et al. analyzed a processed user-interactional part of this data set, consisting of 3488 users and 26404 viewing sessions, to model binge watching behavior for VoD services; we consider the full set of users in our analysis and focus on the transactional data base instead. We present few summary statistics based on the transactional data-base; these preliminary statistics and observations lead to the belief that there is latent structure in the users consumption patterns and guide the pre-processing stage to construct features from raw transaction logs. Note that our results on characterizations of user behavior presented under **Labels in user personas** can be viewed as more detailed and refined summaries. The transactions break up into two types- rentals and purchases with **88**% rentals and **12**%. The price categories of rentals vary from $\mathbf{0-5}$\$ with higher price categories falling in the $\mathbf{3-5}$\$ range. The purchases range as high as **25**\$, mainly for new movies and tv series. The purchases greater than **10**\$ in value are considered as higher end transactions. Most transactions occur in the lower price categories of both types of transactions with only **10**% of consumers transacting in the higher price ranges. A transactional perspective of the content catalogue is observed through segregation of transactions into **15**% TV shows and **85**% movies, with the movie *Frozen* being the most consumed content in the catalogue. The dominant genres in the transactions are *Drama (***18**%*), Comedy (***10**%*), Action (***10**%*), Family (***9**%*), Animation (***7**%*), Thriller (***6**%*), Biography (***5**%*), Sci-Fi (***4**%*), Crime (***4**%*)* etc, with the crucial observation that while some users (**20**%) tend to prefer more family-friendly content (Family, Animation, Super-hero). Other segments (**80**%) consume genres such as drama, horror, comedy etc. The time of transactions is seen mostly to range between evenings and nights, evenly split between weekdays (Monday-Friday) and weekends. As part of the pre-processing of raw logs, barely active users (spend less than **1** dollar in a certain month of activity) and one-time deal hunters (transact only once and never return) were filtered out to prevent cluster centers being pulled to **0**. Summarized information of this data is uncovered from latent structure as cluster centers and cluster sizes, which preserves anonymity of individual users.

Construction of user personas

We offer a behavioral segmentation based on spending traits, content preferences and transactional habits of users with interest in above characterizations stemming from intuition, domain knowledge and business goals. Before, we delve into the model for describing these characterizations, we discuss the timeline and granularity of comparison and a brief description of features (constructed as aggregates over a month of transactions, binned appropriately), crucial in

excavating meaningful latent structure in data.

Timeline of comparison

Transaction logs consist of time-series data. To uncover consequential latency, we make a careful choice as to how the timelines of different users are compared with regard to the following 2 aspects:

Temporal alignment of user timelines: User timelines can be aligned on a calendar basis or on a tenure basis. In the calendar basis, transactions of different users happening at the same calendar dates, for instance in January 2014, are compared against each other. Aligning timelines according to a calendar basis allows to detect seasonalities(holidays, end-of-year movie releases), and effects of specific events happening at a particular date (TV-show new episode/season release or end). On the other hand, in the tenure basis, the first transaction of a given user defines the birth of the user timeline, and transactions of different users are compared when they happen at the same age of the user in the system. For instance, if user A made his first transaction on January, 15th 2014 and user B made his first transaction on April, 10th 2014, comparisons would be drawn for their first month of transactions between Jan. 15th-Feb 14th 2014 for user A and April 10th-May 9th for user B. Aligning user timelines on a tenure basis allows observations on how users age in the system, may help in understanding behavioral phases and in predicting churn.

Temporal granularity: Timestamps in transaction logs can be specified up to seconds or even milliseconds. When building features based on time-series, the question arises as to the granularity at which events should be grouped to devise the desired features. Transactions can be aggregated at a monthly/ weekly/ daily/hourly granularity. For instance, to compute a count feature at the monthly granularity, transactions happening within the same 30 day period will be aggregated. The granularity level affects the detection of behavioral patterns and cycles.
In this work, user timelines are aligned on a **tenure basis**, and events are considered at a **monthly (30 days) granularity**. The first transaction of a user marks the beginning of its timeline, and user's transaction history is divided into successive periods of 30 days each. Our choice of a monthly granularity is guided by elementary analysis of the transaction logs which show unstable structures with weekly granularity– *Weekly logs were too short a period to capture behavioral patterns*–, and a flat structure at a quarterly granularity–*Quarterly logs were too long to capture the dynamism in user labels due to an over-cumulation effect of data*. Our choice of a tenure basis was motivated by the business need to understand the evolution and phases of user behavior along their transaction histories; this helps model their dynamic behavior, predict loss of interest in system, predict lifetimes etc.

Feature engineering

We next give a description of features that we build from raw logs to unearth latent structure. These features are aggregates of transactions at monthly granularity, binned into categories. The choices of binning, arising from a combination of knowledge from histograms depicting the data and domain information, lead to the below features.

Monthly Expenditure (ME) characterizes spending behavior: Each feature is the total net amount spent in one month by a user in a given transaction type (rental/purchase) and a given price category (5 categories for rentals, 8 categories for purchases).

Transaction frequency (TF) characterizes economic behavior: Features are counts of transactions binned into 2 price categories in rentals and 4 price categories in purchases.

Dominant genres (DG) indicates content preference: Features are monthly counts of a user transactions in the 15 most popular genres: *Drama, Comedy, Action, Family, Animation, Thriller, Biography, Sci-Fi, Crime, Super Hero, Comedy-Drama, Fantasy, Horror, Romance, Kids, Miscellaneous.* Note that if we use lower dimensional linear combinations of genres as features obtained through a PCA (Principal Component Analysis) of genre input feature space, we lose clustering structure in addition to interpretability. Rather, binning transactions into dominant genres helps uncover interesting latent structure that provides interpretable user characterizations.

Content recency (CR) indicates freshness preference: Features are counts binned into ranges of content release year: *Old:* < 1990, *Nostalgia:* $1990-2000$, *Not New:* $2000-2010$, *Recent:*$2010-2013$ *and Latest:* $2014-2015$.

Time & day of transaction (TDT) reflects transactional habits: We process time-stamps of transactions to generate the day of the week and time of transaction as per the geographic region of the user. We use counts binned into weekdays or weekends and 4 time slots: *10 AM-5 PM (Office Hours), 5PM-10 PM (evening and night), 10PM-5AM (late night).* No transactions are made in morning pre-office hours.

A mixture model for latent characteristics

To fix notations for this section, we have a $n \times d$ feature matrix

$$X^T = (x_1, x_2, \cdots, x_n),$$

with $x_i \in \mathbb{R}^d$ representing the feature vector of user i in a sample of n users. We propose a parametric approach, a mixed multinomial model **MMM**, implementations of which are seen in Dunson (2000); Chiu et al. (2001); Vermunt and Magidson (2002); Hedeker (2003); Erosheva et al. (2004) to describe user labels based on count data. The choice of a multinomial distribution is motivated from the very fact that it is an extremely natural model for count feature vectors. The iterative EM algorithm applied to estimate the mixing proportions and

the parameters in mixed multinomial distribution, is in itself a very powerful mechanism, with one of its many merits being the ability to deal with missing features. A more common model is the Gaussian mixture model GMM, applied in Reynolds and Rose (1995); Permuter et al. (2003).
A MMM is based on the assumption that the rows of X are independent draws from a multinomial model, that is $x_i \sim MN(d, \theta_{Z_i})$ where Z_i being a latent variable taking values $j \in [K]$, K- representing the number of clusters; independent of X_i. We have a hierarchical structured model as

$$Z_i \overset{iid}{\sim} MN(1, \pi);\ X_i | Z_i \overset{ind}{=} z \sim MN(d, \theta_z),$$

with $\pi = (\pi_1, \cdots, \pi_K)$ representing mixing probabilities for the K clusters and $\theta_j = (\theta_{j,1}, \cdots, \theta_{j,d});\ j \in [K]$ respresenting the multinomial parameters in density given latent factor $Z = j$.

The mixing probabilities π and parameters of mixture model $\theta_j, j \in [K]$ are estimated using an EM algorithm as proposed in Dempster et al. (1977). We outline the E and M steps for the (t)-th iteration of the algorithm for the MMM based on iterates $\pi^{(t)}$ and $\theta^{(t)}$–
E-step: computes the posterior probabilities given estimates of parameters π and θ_{Z_j} of the t-th iteration, that is

$$\begin{aligned}\tau_{i,z}^{(t)} &= \mathbb{P}(Z_i = z | x_i; \pi^{(t)}, \theta^{(t)}) \\ &= \frac{\mathbb{P}(X = x_i | Z_i = z;\ \theta^{(t)}) \pi_z^{(t)}}{\sum_{j=1}^{k} \mathbb{P}(X = x_i | Z_i = j; \theta^{(t)}) \pi_j^{(t)}}.\end{aligned}$$

where $\mathbb{P}(X = x_i | Z_i = j; \theta^{(t)}) = \Pi_{v=1}^{d} \theta_{j,d}^{x_{i,d}}$.
M-step: maximizes the Expected Complete Log Likelihood (ECLL) to refine estimates of parameters π and θ_j for $j \in [K]$;

$$\theta^{(t+1)}, \pi^{(t+1)} = \arg\max \mathbb{E}(\mathcal{L}(\theta^{(t)}, \pi^{(t)}; X))$$

with ECLL $\mathbb{E}(\mathcal{L}(\theta^{(t)}, \pi^{(t)}; X))$ given by–

$$\begin{aligned}&\sum_{i=1}^{n} \sum_{j=1}^{k} \tau_{i,j}^{(t)} \times \log(\pi_j \cdot \mathbb{P}(X = x_i | Z_i = j; \theta)) \\ &= \sum_{i=1}^{n} \sum_{j=1}^{k} \tau_{i,j}^{(t)} \log \pi_j + \sum_{i=1}^{n} \sum_{j=1}^{k} \sum_{v=1}^{d} x_{i,d} \tau_{i,j}^{(t)} \log \theta_{j,d}\end{aligned}$$

yielding estimates

$$\pi_j^{(t+1)} = \frac{1}{n} \sum_{i=1}^{n} \tau_{i,j}^{(t)},\ \theta_{j,v}^{(t+1)} = \sum_{i=1}^{n} x_{i,v} \tau_{i,j}^{(t)} / d \sum_{i=1}^{n} \tau_{i,j}^{(t)}.$$

Hard clustering assignments are obtained by calculating $S_{x_i} = \arg\max_z \mathbb{P}(Z_j = z | x_i; \pi, \theta)$, with $S_{x_i} \in [K]$ for $i \in [n]$.

As is evident from the model above, we need to decide the number of clusters K before we can fit the above latent factor model. Emphasizing on interpretability of labels along with stability and representativeness of clusters, we use a combination of -**stability**, **dominance** and **interpretability** (as described in the appendix) instead of a naive $WCSS$-criterion to evaluate our clustering performance and decide on the number of clusters.

User personas and insights

Having described our approach and methods of constructing personas, we present the summaries of our user personas based on preferential and behavioral pattern. As highlighted earlier, our emphasis goes beyond discovering latent structure. The significant highlights of these labels are clear characterizations of behavior of consumers assigned to these labels. Further, we give other interesting insights which can lead to future business actions to understand evolving patterns of dominant and niche behavioral characterizations.

User labels for behavioral characterizations

We describe the latent structure, excavated from the binned features that were constructed from raw logs. This is supplemented with detailed tables for cluster centers and sizes in the appendix. To avoid obscurity in numbers, we enlist the labels with the cluster sizes reported in percentages (beside the label) and for each label, we describe the characterization with a brief explication of consumer behavior in that bucket. Note that all the labels have clear characterizations lending interpretations for business plans.
Monthly Expenditure: Cluster centers represent monthly expense in each of the 13 price categories (5 rental price categories and 8 purchase categories).

- **Economic Renters** (71%) : 10$ in all in a month of activity, with smaller amounts of $2-3$$ each in the higher renting price categories
- **Heavy Renters** (21%) : 17$ in total, mostly around 13$ in the $3-5$ rental price category
- **Movie Buyers** (4.5%): 32$ in total with one purchase on average in the $16-20$ price category and one fourth of monthly expenses in higher rentals and lower priced purchases.
- **Movie Buffs** (2.5%): 60$ in all, with around 3 purchases in $10-16$ and around 7 dollars in $16-20$ price categories.

Frequency of Transaction: Cluster centers denote transaction counts in 6 price ranges, the labels uncovered are–

- **Frequent High-End Renters** (61%): over 85% transactions in rentals above 3$
- **Frequent Low-End Renters** (21%): over 60% and 30% transactions in rentals below and above 3$ respectively
- **Frequent Movie Buyers & Sporadic Renters** (12%): 45% transactions in purchases in $8-16$$ with about 35% transactions in rentals too.

- **Frequent Low End Purchasers** (6%): 80% transactions mostly in the $0-8$ price category of purchases

Dominant Genre of Content Consumed: The three prime clusters recovered with cluster centers being percentage of monthly transactions in 16 genres are–

- **Happy Family** (23%): content qualifying as family watch with distribution being family genre (28%)–the most consumed genre, followed by animation (20%), comedy (13%); but no/ almost no crime, horror, romance, thriller
- **Drama-Comedy:** (40%) content with dominant genres– drama (28%), followed by biography (10%), comedy (10%), bit of romance but little or almost nothing as compared to other clusters in terms of consuming family, horror, action,crime
- **Action-Horror-Thrill:** (37%) dominant genre is action (20%), followed by drama (15%), thriller (12%), sci-fi (8%), comedy (6%), horror (5%), but little or almost nothing as compared to other clusters in terms of consuming family, comedy-drama, fantasy.

Recency of Content consumed: We cluster the 3 genre clusters individually based on the count matrices binned as per release year of content to observe characterizations for recency of content. Denoting cluster sizes when clustered within the family cluster as F and within the non-family clusters (Drama-Comedy,Action-Horror-Thriller) as NF, labels are–

- **Latest** (F-35%, NF-45%): 85% of transactions with release year 2014-15.
- **Recent** (F-35%, NF-40%) : 85% of transactions with release in 2010-13.
- **Nostalgic** (F-30%, NF-15%): About 30% with release in 2000-09 followed by recent and latest content in the remaining 65% of transactions

Time & Day of Transaction: Based on habit or preference to transact at a certain time and day of the week, clusters with centers representing counts in each time category of weekday/ weekend are–

- **Weekend Evening & Night** (24%): 65% of transactions on weekend nights, followed by 25% in evening.
- **Weekday Evening & Night** (24%): 70% of transactions on weekday nights, followed by 20% in evening.
- **Weekend & Weekday Night** (42%): 45% and 35% of transactions on nights of weekdays and weekends
- **Weekend Day & Night:** (10%): 25% of transactions on weekend day time and 60% in weekend nights

Insights into labels

The user personas stated above lend insights that validate common intuitions about consumer behavior and thereby, can render easy motivations for actionable schemes. We describe temporal evolution of user labels - "how dominant characterizations stay stable while niche ones undergo change". The other attractive property of the discovered latency is stability on a population level. We show a divisive structure in the labels as we increase the number of clusters; this gives us insights into layers within a single characterization. Further, we also

use our approach to discover inter-relations between different characterizations as layered clustering instead of a joint one. While the former that discovers layers within a single characterization is not imposed by our algorithm, the latter uses it to discover a layer of characterization within another.

Temporal nature of labels - *"stability of macro characteristics":* An interesting and important aspect of the personas constructed herein is that the clusters uncovered stay stable in terms of size and composition on a population level. Such an observation suggests consistency of population characteristics along tenure timelines. Hence, the reported labels hold valid on a macro-level at any time point of tenure evolution. Our personas do have this attractive property of consistency which allows us to use the same model for population characterizations. At the same time, they also succeed in explaining individual dynamism. That is, migrations do happen on a user level, which means that individual user labels are not static. Our results show these migrations in most categories are never drastic in nature; while we do observe few interesting migrations into far-off labels. These migrations can be explained as *dominant characterizations reflecting spending capacity, content preferences and habits* stay stable over time and are not prone to change; while *niche characterizations* are more prone to change.

As specific examples, often we see users migrating from higher price end categories to the lower ones, as they come back to transact in the higher end with only availability of new products of their interest. Such users represent the niche segment with economic capacity to transact in the higher price categories. However, the dominant segment represented by economic renters stays stable across tenure timeline, migrating to only to clusters with closer centers. Another niche segment is a proportion of people who buy content in the "happy family" label; they move to other other labels of genre consumption as they also buy content for individual consumption that's different from content consumed with family.

To summarize, our user personas suggest that user preferences might change at the micro level, but macro tendencies are seen to exhibit stability; thereby, our personas succeed in capturing the overall inclination and tendencies of consumers, across dominant and the migrations of the niche segments of population, along tenure timelines. The below pictorial representation shows evolution of user labels along tenure timelines of two users, the first row corresponds a member of niche category and the second row corresponding to a member of the dominant category is shown.

Natural hierarchical structure of clusters We see a very natural, divisive, hierarchical structure (not imposed through algorithm) to our perosnas, as we increase number of clusters. This lends interesting interpretations on the sub-population of consumers within broad segments. Such a structure is not imposed by our mixture model, but inherently displayed by the data. We highlight some examples to highlight the insights from this observation. Upon clustering consumers based on monthly expenditure into two clusters, cluster

Figure 1: Temporal nature of labels: dominant spending tendency in evolution of labels in second row stays stable (migrations only to nearby clusters) while niche tendency of buying in the higher priced categories, depicted in the first row, is more prone to change

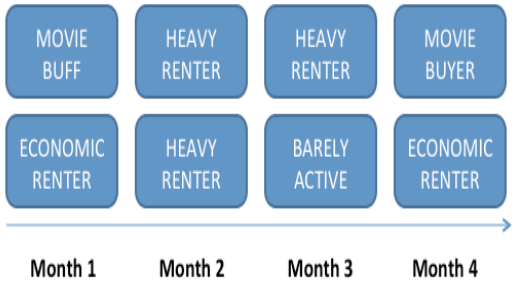

centers represent renters and purchasers. Renters break up into economic and heavy renters with 3 clusters. Purchasers mostly decompose into two clusters– movie buyers and movie buffs with 4 buckets. Upon clustering count data representing dominant genres consumed by users into two clusters, we see a segment preferring family content over content not qualifying predominantly as family watch. With three clusters, non family content consumers decompose into two buckets- one that consumes drama, comedy etc while other prefers thrill inducing content. Figures 2 and 3 show these examples pictorially.

Figure 2: Hierarchy of spend labels: this depicts that the first two clusters are that of renters and purchasers, which later decompose into the smaller segments within each label.

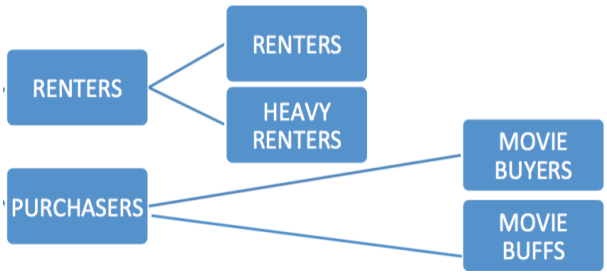

Layered structure of clusters: We also have insights from performing a layered clustering model, i.e., looking for labels for a characterization like genre preference within clusters for spending behavior. We do this in place of a joint clustering of features to avoid loss of meaningful clustering structure. For instance, clustering based on genre preference within the clusters characterizing economic behavior is similar across all economic clusters. Similarly the clusters for spending behavior are similar across different genres. This observation validates that genre preferences of consumers is independent of their

Figure 3: Hierarchy of genre labels: The two major labels upon setting the number of clusters to 2 is that of the users consuming and not consuming family content. Further clustering splits the consumers that do not transact on family content typically into the two labels: that prefer comedy-drama-romance and action-thriller-horror

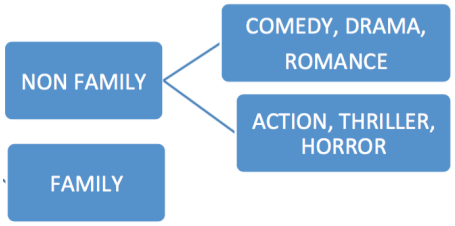

economic budget. A similar observation goes for recency and economic behavior. On the other hand, we see different clustering results for recency of content when clustered within the genre clusters, with the category preferring family content showing more inclination towards more classic content that the other drama-based or thrill inducing categories that prefer more recent content. Layered clustering within constructed personas also helps relate different consumer characteristics. Figure 4 depicts the above discussed illustrations.

Figure 4: Layers between characterizations: same labels for genre preference within each of the spending labels, but different distributions for recency preference within the genre preference labels

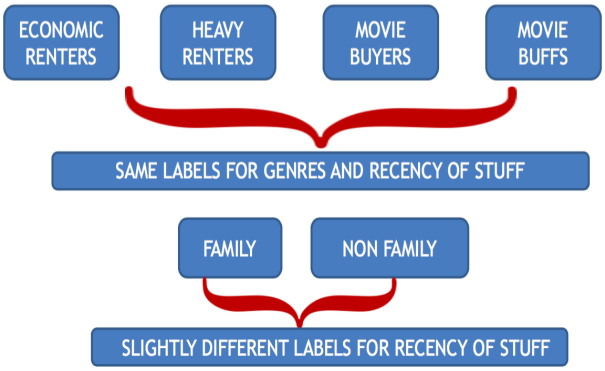

Integration of personas in personalization

We devote this section to discuss effective integration of information from personas into the personalization. Specifically, we focus on a CTR (Click to rate) predictive model where we predict $p_{u,i}$, the probability that an user u will transact on an item i. We highlight the efficient use of persona information, either in the form of soft clusters or hard clusters at the training stage based on a novel

set of guarantees, namely **relevance-scalability-interpretability**. Though we show our results in this specific predictive setting, the scope of utilizing information from personas extends to other popular models in collaborative filtering. We conclude the paper by discussing such possibilities, where one can integrate personas into commonly used models and expect to reap similar benefits.

CTR: tradeoff in relevance-scalability -interpretability

We demonstrate the utility of user personas in a CTR predictive model. We use an ℓ_1 penalized logistic regression model trained per item to predict transactional probabilities. Such a finer-grained model at the item level can capture the item's specific attraction for users, resulting in more accurate predictions as advocated in Zhang et al. (2016). Adding to the challenge in such models is the sparsity of the transactional database, with about **1**% users transacting on any given item. For prediction purposes, we sub-sample for every positive sample (users who transacted) 5 negative samples (users who did not transact) to overcome this imbalance and to avoid bias towards the outcome "not transacting at all".

Typically, the standard measure to assess the performance of such predictive models is by using an evaluation metric $\mathcal{F}$ on a test set, denoted by $\mathcal{F}(\widehat{\text{label}}_{M^*}, \text{label}_{\text{test}})$, with the training model M^* giving predicted labels $\widehat{\text{label}}_{M^*}$. In the current work, we offer a balance between scalability of the training model and interpretability of features alongside maintenance of relevance of predictions; a more wholesome and practical assessment of model performance. We define the merit of a trade-off between these criteria:-

Relevance :-deliver relevant recommendations to users, quantified by the quality of prediction in transactional probabilities or amounts; we use the mean AUC over the **100** most popular items in the content catalogue as a measure of predictive performance

Scalability :-reduce size of input feature and sample space (leads to reduction in regression dimension) by using lower dimensional summary features or by training a model clusterwise respectively. We advocate use of persona information encoded as soft clustering features or incorporated as hard clustering information via a model trained per cluster. This brings significant reduction in regression dimensions which in turn, facilitates storage and future use of feature vectors in other predictive models

Interpretability :-retain meaning to feature space as opposed to random lower dimensional projections which seldom lend business insights and are limited in their scope of extension to other models. With a meaningful feature set, we can expect reutilizing the same features in a host of predictive tasks

and in easy debugging of models in the future. While relevance and scalability can be quantified, there is no measure that quantifies interpretability. We use a subjective judgement here; our personas have clear characterizations.

The trade-off in these criteria arises as we can use the count features (that we used to recover latent user labels) in a predictive model and expect to to have high predictive power. On the other hand, we gain in scalability through a reduction in regression dimension as highlighted in Table 2 with the use of persona information either as soft/hard clustering information. This gain with the integration of persona information takes place at the cost of losing some predictive ability. In the results that follow, we show that we can achieve significant reductions in regression dimensions and subsequently, in the complexity of fitting a model at a cost of losing a mere 2% in predictive power with soft cluster features or training models per cluster. The baseline of predictive power is the model using the count features as the predictors. Also, we retain clear meaning of features which offer clear characterizations of consumer behavior. This justifies the term "trade-off". We explain this in more details below. To summarize, we advocate the use of persona features as interpretable, lower dimensional vectors that preserve predictive power of models. An added advantage of incorporating these summary features in a model with sub-sampled users is preservation of privacy of individual users. We also preserve anonymity of individual transactions in using summaries over a random set of users.

To describe our model and results, we denote X_u as the feature vector corresponding to user u containing information from 3 characterizations: ME (monthly expenses), DG (dominant genre), CR (content recency). Information from personas can be used in X_u in different ways, yielding different feature spaces and consequently, different models. In particular, we construct feature vectors using the personas on ME, DG and CR in the following different forms (denoted by **(c), (s), (h)** and **(-)** respectively:

(c) a feature vector with distribution of ME in price categories and/ or count vectors for DG/CR in feature bins (directly using the constructed features) : no information from personas

(s) a feature vector of soft clustering values, i.e. distances of the constructed features from the respective cluster centers : using persona information in the form of soft-clusters

(h) incorporate hard clustering information for a characterization by training model cluster-wise : using persona information in the form of hard-clusters

(-) not include any information from a characterization at all : not using persona information at all.

The observations from implementing models based on personas are high-

lighted below: We achieve a **gain in relevance** with information from each added characterization, either in the form of soft clustering/ hard clustering/ count feature. The below table shows the relevance of each characterization in the CTR model. CR (recency) is seen to the most informative characterization adding the most to AUC; this says that the recency of a content item is highly predictive of the consumer's interest in transaction.

Table 1: Mean AUC: most informative feature

Recency	Genre	Economic	AUC
c	c	c	0.85
-	-	c	0.74
-	c	-	0.73
c	-	-	0.79

We reduce **size of regression** at the cost of only 2% loss in predictive performance in terms of mean AUC over items, while retaining interpretability (as opposed to dimensionality reduction in feature space via a random projection onto a lower dimensional space). Denote n_i as the samples per item and $n_{i,c}$ as samples per item, per cluster, p the number of predictive features, O as the complexity of regularized logistic with sample size and regression dimension. The below table compares different models- illustrating how our proposed integration of user personas into personalized recommendation achieves a tradeoff between relevance and scalability (noting that interpretability comes alongside using summary information from personas).

The baseline model in the first row of the table; representing the model with all count features **(c)**. We see a significant reduction in the predictive power when we do not incorporate any information from the recency feature, this is depicted by the fourth row of the table. When we train a model per recency cluster using **(h)**, we lose 1% of predictive power, but reduce the sample size for the training model on each cluster. This leads to a significant reduction in the sample size of the regression model from n_i to $n_{i,c}$. We see a similar predictive power when we use soft clustering recency feature **(s)**, but a significant reduction in the size of feature space. While we do not incorporate all 64 combinations of **(c)**, **(s)**, **(h)**, **(-)**; we see that using soft clustering features for all the 3 characterizations leads to a loss of just 2% AUC. But, the computational gain is significant even in a simple regression model that scales in complexity as p^2 with the size of feature space. This is represented in the last row of below table.

Table 2: AUC: complexity and dimensionality of sample and feature space

Recency	Genre	Economic	AUC	n	p	Complexity
c	c	c	0.85	n_i	140	$O\left(n_i \times 140^2\right)$
s	c	c	0.84	n_i	128	$O(n_i \times 128^2)$
h	c	c	0.84	$n_{i,c}$	116	$O(n_{i,c} \times 116^2)$
-	c	c	0.75	n_i	116	$O(n_i \times 116^2)$
s	s	s	0.82	n_i	40	$O(n_i \times 40^2)$

Persona based collaborative filtering models

We finally discuss few models based on popular collaborative filtering techniques that can incorporate information from personas to retain predictive power while gaining in scalability for practical implementations.

User based nearest neighbor similarity: This approach is based on a similarity metric $\text{sim}(u, v)$ (examples include Jaccard, cosine etc.) to predict a weighted average rating based on similarity between users who transacted on the same items. Denoting $U(i)$ as the set of users who transacted on the same item i, amount of money user u is willing to spend on item i can be predicted as

$$r_{u,i} = \frac{\sum_{v \in U(i)} \text{sim}(u, v) r_{v,i}}{\sum_{v \in U(i)} \text{sim}(u, v)}.$$

and probability that user u transacts on item i as

$$p_{u,i} = \frac{\sum_{v \in U(i)} \text{sim}(u, v)}{\sum_{v} \text{sim}(u, v)},$$

with $1_{v,i}$ as indicator if user v transacted on item i. Similarity approaches have scaling issues with high computational cost associated with searching through set of users or even the top K similar users in the set $U(i)$. Persona information can bring in gain in prediction accuracy, also offering better scalability via limiting search of top K neighbors to already formed personas.

While we could use clusters from most representative time point of activity for predictions, we can also use temporal persona information for prediction with the scope of leveraging differently on time points through a weighted similarity prediction along the tenure timeline. Denoting time points of transaction history (months of tenure timeline) as t with weights w_t (that can be tuned) and features–u_t for user u, $U(t, i)$ as the set of users who transacted on the same item i and $C(t, u)$ the set of users present in the same cluster as user u at time t, ratings at a time point T leveraging on temporal history till time T can be predicted as–

Table 3: Ratings in CF: Clustering buckets $C(u)$

$r_{u,i}(T) = \frac{\sum_{t \le T} \sum_{v \in U(t,i) \cap C(t,i)} w_t \text{sim}(u_t, v_t) r_{v,i}}{\sum_{t \le T} \sum_{v \in U(t,i) \cap C(t,i)} w_t \text{sim}(u_t, v_t)}$
$p_{u,i}(T) = \frac{\sum_{t \le T} \sum_{v \in U(t,i) \cap C(t,i)} w_t \text{sim}(u_t, v_t) r_{v,i}}{\sum_{t \le T} \sum_{v \in C(t,i)} w_t \text{sim}(u_t, v_t)}$

Latent factor model: Without clustering information, the vanilla model with latent factors q_i for item i and p_u for user u is

$$\hat{r}_{u,i} = \mu + b_i + b_u + q_i^T p_u$$

solved either through stochastic gradient descent or alternating least squares. Letting $\mathcal{A}$ to be a set of attributes and a a cluster for $\mathcal{A}$, persona information can be incorporated into the above model by (see below table for ratings in each case in that order)–

1. adjusting for biases per cluster
2. enhancing user representation in the form of latent factors for cluster memberships learnt with $y_a \in \mathcal{A}$–a latent factor for each cluster a in set of characterizations (see Koren et al. (2009))
3. hard wiring clustering information as features in the form of an enhanced user feature with a latent component p_u concatenated with known added features $\tilde{p}_u$
4. training latent factor model per cluster with c_a being clusters corresponding to some attribute a.

Table 4: Ratings in CF: Adding Clustering Info to Vanilla

(a)	$\hat{r}_{u,i} = \mu + b_i + b_u + \sum_{a \in \mathcal{A}(u)} b_a + q_i^T p_u$
(b)	$\hat{r}_{u,i} = \mu + b_i + b_u + q_i^T (p_u + \sum_{a \in \mathcal{A}(u)} y_a)$
(c)	$\hat{r}_{u,i} = \mu + b_i + b_u + \tilde{q}_i^T (p_u : \tilde{p}_u)$
(d)	$\hat{r}^{c_a}_{u,i} = \mu_{c_a} + b_i + \mathbb{I}_{u \in c_a} b_u + \mathbb{I}_{u \in c_a} q_i^T p_u$

Concluding remarks

The major take-away from this work of modeling latent structure from raw transactional data can be seen as an efficient dimensionality reduction in regression size. All of these merits come with maintenance of predictive power and preservation of interpretability for actionable business goals. The temporal personas constructed through a mixture model approach capture behavioral patterns of VoD users as they age in the system. As highlighted, our personas do possess the essential stability on a macro level, while representing the dynamic niche characterizations simultaneously. Hence, our mixture approach together with the choices of granularity and timeline of comparison and the engineered features give rise to a consistent and robust latent model. Any insights derived from a study of personas at any time point are also likely to apply to future clusters. We also illustrated how information from efficiently built personas can balance a much practical and vital *relevance-scalability-interpretability* trade-off in recommendations. User personas can find applications in building other models like the collaborative filtering models stated above for better scalability while preserving predictive power and business insights. Another untapped area of application is churn analysis, see Kapoor et al. (2014), aiming to improve user retention and interest. One can create user buckets based on longevity in system or use existing personas to predict when users slip into a state of coma in

the system. This work suggests methods, guarantees and perspectives that are certainly not limited to the VoD streaming space; such a construction of personas and integration into predictive models can be extended to other domains of personalization.

References

Deepak Agarwal, Bee-Chung Chen, and Pradheep Elango. Spatio-temporal models for estimating click-through rate. In *Proceedings of the 18th international conference on World wide web*, pages 21–30. ACM, 2009.

Rakesh Agrawal, Manish Mehta, John C Shafer, Ramakrishnan Srikant, Andreas Arning, and Toni Bollinger. The quest data mining system. In *KDD*, volume 96, pages 244–249, 1996.

Greg M Allenby and Peter E Rossi. Marketing models of consumer heterogeneity. *Journal of econometrics*, 89(1):57–78, 1998.

Asim Ansari, Skander Essegaier, and Rajeev Kohli. Internet recommendation systems. *Journal of Marketing research*, 37(3):363–375, 2000.

Chidanand Apte, Bing Liu, Edwin PD Pednault, and Padhraic Smyth. Business applications of data mining. *Communications of the ACM*, 45(8):49–53, 2002.

Judy Bayer. Customer segmentation in the telecommunications industry. *Journal of Database Marketing & Customer Strategy Management*, 17(3-4):247–256, 2010.

David Besanko, Jean-Pierre Dubé, and Sachin Gupta. Competitive price discrimination strategies in a vertical channel using aggregate retail data. *Management Science*, 49(9):1121–1138, 2003.

Amit Bhatnagar and Sanjoy Ghose. A latent class segmentation analysis of e-shoppers. *Journal of Business Research*, 57(7):758–767, 2004.

Derrick S Boone and Michelle Roehm. Retail segmentation using artificial neural networks. *International journal of research in marketing*, 19(3):287–301, 2002.

Deepayan Chakrabarti, Ravi Kumar, and Andrew Tomkins. Evolutionary clustering. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 554–560. ACM, 2006.

Olivier Chapelle, Eren Manavoglu, and Romer Rosales. Simple and scalable response prediction for display advertising. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 5(4):61, 2015.

Mu-Chen Chen, Ai-Lun Chiu, and Hsu-Hwa Chang. Mining changes in customer behavior in retail marketing. *Expert Systems with Applications*, 28(4):773–781, 2005.

Tom Chiu, DongPing Fang, John Chen, Yao Wang, and Christopher Jeris. A robust and scalable clustering algorithm for mixed type attributes in large database environment. In *Proceedings of the seventh ACM SIGKDD international conference on knowledge discovery and data mining*, pages 263–268. ACM, 2001.

Florence Corpet. Multiple sequence alignment with hierarchical clustering. *Nucleic acids research*, 16(22):10881–10890, 1988.

Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society. Series B (methodological)*, pages 1–38, 1977.

José G Dias and Jeroen K Vermunt. Latent class modeling of website users? search patterns: Implications for online market segmentation. *Journal of Retailing and Consumer Services*, 14(6):359–368, 2007.

David B Dunson. Bayesian latent variable models for clustered mixed outcomes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 62(2):355–366, 2000.

Michael D Ekstrand, John T Riedl, and Joseph A Konstan. Collaborative filtering recommender systems. *Foundations and Trends in Human-Computer Interaction*, 4(2):81–173, 2011.

Elena Erosheva, Stephen Fienberg, and John Lafferty. Mixed-membership models of scientific publications. *Proceedings of the National Academy of Sciences*, 101(suppl 1):5220–5227, 2004.

Nir Friedman and Stuart Russell. Image segmentation in video sequences: A probabilistic approach. In *Proceedings of the Thirteenth conference on Uncertainty in artificial intelligence*, pages 175–181. Morgan Kaufmann Publishers Inc., 1997.

Johanna Gummerus, Veronica Liljander, Minna Pura, and Allard Van Riel. Customer loyalty to content-based web sites: the case of an online healthcare service. *Journal of services Marketing*, 18(3):175–186, 2004.

Donald Hedeker. A mixed-effects multinomial logistic regression model. *Statistics in medicine*, 22(9):1433–1446, 2003.

Yifan Hu, Yehuda Koren, and Chris Volinsky. Collaborative filtering for implicit feedback datasets. In *2008 Eighth IEEE International Conference on Data Mining*, pages 263–272. Ieee, 2008.

Anil K Jain, Robert P. W. Duin, and Jianchang Mao. Statistical pattern recognition: A review. *IEEE Transactions on pattern analysis and machine intelligence*, 22(1):4–37, 2000.

Kamel Jedidi, Harsharanjeet S Jagpal, and Wayne S DeSarbo. Finite-mixture structural equation models for response-based segmentation and unobserved heterogeneity. *Marketing Science*, 16(1):39–59, 1997.

Tianyi Jiang and Alexander Tuzhilin. Segmenting customers from population to individuals: Does 1-to-1 keep your customers forever? *IEEE Transactions on Knowledge and Data Engineering*, 18(10):1297–1311, 2006.

Tianyi Jiang and Alexander Tuzhilin. Improving personalization solutions through optimal segmentation of customer bases. *IEEE transactions on knowledge and data engineering*, 21(3):305–320, 2009.

Stephen C Johnson. Hierarchical clustering schemes. *Psychometrika*, 32(3): 241–254, 1967.

Komal Kapoor, Mingxuan Sun, Jaideep Srivastava, and Tao Ye. A hazard based approach to user return time prediction. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1719–1728. ACM, 2014.

Yehuda Koren, Robert Bell, Chris Volinsky, et al. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009.

Gueorgi Kossinets and Duncan J Watts. Empirical analysis of an evolving social network. *science*, 311(5757):88–90, 2006.

Wei Li, Xuemei Wu, Yayun Sun, and Quanju Zhang. Credit card customer segmentation and target marketing based on data mining. In *Computational Intelligence and Security (CIS), 2010 International Conference on*, pages 73–76. IEEE, 2010.

Geoffrey McLachlan and Thriyambakam Krishnan. *The EM algorithm and extensions*, volume 382. John Wiley & Sons, 2007.

Geoffrey McLachlan and David Peel. *Finite mixture models*. John Wiley & Sons, 2004.

Paul D McNicholas and Thomas Brendan Murphy. Model-based clustering of microarray expression data via latent gaussian mixture models. *Bioinformatics*, 26(21):2705–2712, 2010.

Michael J Pazzani and Daniel Billsus. Content-based recommendation systems. In *The adaptive web*, pages 325–341. Springer, 2007.

H Permuter, J Francos, and Ian H Jermyn. Gaussian mixture models of texture and colour for image database retrieval. In *Acoustics, Speech, and Signal Processing, 2003. Proceedings.(ICASSP'03). 2003 IEEE International Conference on*, volume 3, pages III–569. IEEE, 2003.

Douglas Reynolds. Gaussian mixture models. *Encyclopedia of biometrics*, pages 827–832, 2015.

Douglas A Reynolds and Richard C Rose. Robust text-independent speaker identification using gaussian mixture speaker models. *IEEE transactions on speech and audio processing*, 3(1):72–83, 1995.

Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th international conference on World Wide Web*, pages 285–295. ACM, 2001.

J Ben Schafer, Dan Frankowski, Jon Herlocker, and Shilad Sen. Collaborative filtering recommender systems. In *The adaptive web*, pages 291–324. Springer, 2007.

Wendell R Smith. Product differentiation and market segmentation as alternative marketing strategies. *Journal of marketing*, 21(1):3–8, 1956.

William Trouleau, Azin Ashkan, Weicong Ding, and Brian Eriksson. Just one more: Modeling binge watching behavior.

A Vellido, PJG Lisboa, and K Meehan. Segmentation of the on-line shopping market using neural networks. *Expert systems with applications*, 17(4):303–314, 1999.

Jeroen K Vermunt and Jay Magidson. Latent class cluster analysis. *Applied latent class analysis*, 11:89–106, 2002.

Michel Wedel and Wagner A Kamakura. *Market segmentation: Conceptual and methodological foundations*, volume 8. Springer Science & Business Media, 2012.

Jing Wu and Zheng Lin. Research on customer segmentation model by clustering. In *Proceedings of the 7th international conference on Electronic commerce*, pages 316–318. ACM, 2005.

XianXing Zhang, Yitong Zhou, Yiming Ma, Bee-Chung Chen, Liang Zhang, and Deepak Agarwal. Glmix: Generalized linear mixed models for large-scale response prediction. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 363–372. ACM, 2016.

Appendix

Evaluation of clustering performance

The clusters in the Results section above satisfy 3 criteria–**stability**, **dominance** and **interpretability** as defined below. We define these below:
$\epsilon - \delta$ **Stable clusters:** Theoretically, clustering stability is achieved if similar results are obtained when applied to several data sets from the same underlying model or of the same data generating process. In our case, clusters are defined to be $\epsilon - \delta$ stable on a population level if the cluster centers of the most similar clusters are ϵ close in ℓ_2 norm and the corresponding cluster sizes vary by a fraction of at most δ when clustering algorithm is repeatedly applied to random sub-samples of size 50% (without replacement) of the original data set.
$K_{\mathbf{max}} - \kappa$**-thresholded Dominant clusters:** Typically, clustering is said to be dominant if clusters obtained are representative of large segments and also, of niche segments of the population. A precise way of obtaining such clusters which capture both the dominant and niche characterizations of a population in our case is by setting the cluster sizes to be above a threshold. In our case, clusters are said to be $K_{\max} - \kappa$-thresholded dominant clusters if cluster sizes are at least greater equal to $\kappa > 0$ and number of clusters is upper bounded by $K_{\max}$.
Interpretable clusters: Interpretability to hard to made precise mathematically. We define it loosely in terms of clear meanings to clusters that can help characterize customers distinctly. Clusters are deemed interpretable if the summarized information from the clusters can be lent to intuitive explanation and can lead to actionable insights that can help improve key business performance indicators.

Clusters: centers and sizes

Below, we give details of clusters of the characterizations: *Monthly Expenditure*, *Frequency of Transaction*, *Dominant Genre of Content Consumed*, *Recency of Content consumed*, *Time & Day of Transaction*. We denote the clusters as Cl-i, $i = 1, 2, 3, 4$. The percentages in brackets in the first column denote the cluster sizes in terms of percentages. For *Monthly Expenditure*, the cluster centers denote the average monthly expenses in each price category. For all other characterizations, the cluster centers show the percentage of total monthly transactions in a certain category.

Table 5: Monthly expenditure: clusters in terms of rental price categories

Rental	0	$0-1$	$1-3$	$3-5$	>5
Cl-1 (71 percent)	0	0	0.97	2.38	2.41
Cl-2 (21 percent)	0	0	1.48	13.02	1.79
Cl-3 (4.5 percent)	0	0	0.82	3.35	2.46
Cl-4 (2.5 percent)	0	0.01	1.79	6.02	3.19

Table 6: Monthly expenditure: clusters in terms of purchase price categories

Purchase	0	$0-3$	$3-5$	$5-8$	$8-10$	$10-16$	$16-20$	>20
Cl-1 (71 percent)	0	0.56	0.01	0.04	0.33	1.07	0	0.44
Cl-2 (21 percent)	0	0.36	0.01	0.03	0.34	0.87	0.04	0.16
Cl-3 (4.5 percent)	0	0.56	0.02	0.10	1.10	3.74	23.95	0.09
Cl-4 (2.5 percent)	0	1.11	0.08	0.39	4.29	39.86	6.49	2.12

Table 7: Frequency of Transaction: rental and purchase price categories

Clusters	$0-3$	>3	$0-8$	$8-16$	$16-20$	>20
-	Ren	Ren	Pur	Pur	Pur	Pur
Cl-1 (61 percent)	10	86	0	3	1	0
Cl-2 (21 percent)	60	33	2	5	0	0
Cl-3 (12 percent)	8	29	4	43	13	3
Cl-4 (6 percent)	5	12	78	4	0	1

Table 8: Dominant Genre of Content Consumed

Cl	Drama	Comedy	Action	Family	Ani	Thrill	Biog	Scifi
Cl-1 (23 percent)	5	13	5	28	20	2	2	2.5
Cl-2 (39-42 percent)	28	10	4	1	0	3	10	0
Cl-3 (38-35 percent)	15	6	20	1	0	12	2	8

Table 9: Dominant Genre of Content Consumed

Cl	Crime	Sup.her	Com-dr	Fan	Horr	Rom	Kids	Misc.
Cl-1 (23 percent)	0	4	2	4	0	0	3	7
Cl-2 (39-42 percent)	4	0	5	7	0	4	0	20
Cl-3 (38-35 percent)	4	4	1	1	5	2	0	14

Table 10: Recency of Content: under cluster consuming family content

Cl	<1990	1990-2000	2000-10	10,11,12,13	14,15
Cl-1 (25 percent)	7	10	35	27	20
Cl-2 (37 percent)	1	1	7	78	13
Cl-3 (38 percent)	1	1	7	13	78

Table 11: Recency of Content: under cluster consuming non-family content

Cl	<1990	1990-2000	2000-10	10,11,12,13	14,15
Cl-1 (16 percent)	8	9	28	32	33
Cl-2 (39 percent)	0	0	3	88	9
Cl-3 (45 percent)	0	0	3	9	88

Table 12: Time & Day of Transaction

Cl	Wk 10-5	Wk 5-10	Wk 10-5	Wd 10-5	Wd 5-10	Wd 10-5
Cl-1 (24 percent)	0	1	10	1	24	62
Cl-2 (42 percent)	4	9	45	3	6	31
Cl-3 (24 percent)	8	18	68	0	0	5
Cl-4 (10 percent)	2	2	10	24	7	56